

Generalization bound for kernel similarity learning

Michael Rabadi
 Spotify
 New York, NY 10011
 mrabadi@spotify.com

October 24, 2016

Abstract

Similarity learning has received a large amount of interest and is an important tool for many scientific and industrial applications. In this framework, we wish to infer the distance (similarity) between points with respect to an arbitrary distance function d . Here, we formulate the problem as a regression from a feature space \mathcal{X} to an arbitrary vector space \mathcal{Y} , where the Euclidean distance is proportional to d . We then give Rademacher complexity bounds on the generalization error. We find that with high probability, the complexity is bounded by the maximum of the radius of \mathcal{X} and the radius of \mathcal{Y} .

1 Introduction

There are many applications where scientists wish to infer similarity between samples. In this framework, we are given a sample $S = (x_1, x_2, \dots, x_m)$, $x \in \mathbb{R}^N$ of points with associated features in some vector space $\mathcal{X} \subset \mathbb{R}^N$ and a corresponding distance matrix \mathbf{D} , where $\mathbf{D}_{ij} = d(x_i, x_j)$, for some arbitrary distance function (similarity) d . The goal is to learn a mapping $h : \mathcal{X} \rightarrow \mathcal{Y}$, where the Euclidean distance in \mathcal{Y} is close to the corresponding distances of \mathbf{D} . Of course it is not sufficient to simply reconstruct \mathbf{D} , since the goal is to infer the distance between new points. Indeed, this framework is completely general and immediately relevant to many scientific and industrial domains.

One notable application comes from neuroscience, where scientists may simultaneously record neural signals and behavioral measurements. One concrete example comes from object recognition research, where a researcher may want to relate behavior (confusion rate between pairs of objects) and neural activity for different objects [3, 4, 6]. In many cases, neuroscientists

will use a linear classifier to infer separability of the neural signal, which implies that a decision making part of the brain could linearly categorize stimuli based on those neural signals (see [3] for review). Unfortunately, this approach does not capture the full behavioral profile and reduces object recognition to a binary behavior. Another approach is to use dimensionality reduction algorithms to infer categorical structure in the neural signals and relate this to behavior [5, 7]. Unfortunately, this approach does not generalize to new data, since all of the pairwise neural comparisons and pairwise behavioral comparisons must be measured. Thus, introducing a new object requires comparing to all m previously measured objects. In many cases, this is impractical. Similarity learning is a natural solution to all of these drawbacks. However, we must first transform a behavioral confusion matrix into a distance matrix. We can do this simply with the following equation, $\mathbf{D}_{ij} = 1 - \mathbf{C}_{ij}$, where \mathbf{C} is the confusion matrix. Thus, \mathbf{D} implies a bounded space \mathcal{Y} with radius $\beta = 0.5$. Indeed, learning the distances between points gives a finer measure of separation in neural signals than simply demonstrating the ability for a binary classifier to classify objects based on the corresponding neural signal. Furthermore, we can infer distances between new objects based on their neural signature and guarantee the total error with high probability. This allows us to precisely decide the number of objects to compare such that we can infer the rest of the behavior with the learned hypothesis h .

Another important example comes from recommendation systems. Suppose, for example, that we have a global encoding of products in some vector space \mathcal{X} , where the distance between two products implies some relationship. We then want to recommend new products to a customer based on their purchase history S . One approach is to simply use the distances between the products in \mathcal{X} , but there is no reason to believe that global similarity between products is meaningful for a customer. Instead, we might want to construct a new space \mathcal{Y} where the distance between products is proportional to a user's tastes. In this case, we might construct a similarity measure d that is related to a user's purchase or browsing behavior. We can then recommend products by selecting new ones that are similar to a user's purchase history in \mathcal{Y} .

Similarity learning is a general framework with diverse applications, however there has been little work on the generalization analysis of the problem. Previously, [2] derived a generalization bound for the case where \mathcal{Y} is a linear transformation of \mathcal{X} . Their bound was in terms of the norm of the transformation matrix \mathbf{A} and a term related to the maximum distance between two points in \mathcal{X} . They also gave explicit bounds for common norms.

Unfortunately their analysis was limited to linear transformations of \mathcal{X} .

Here, we give new Rademacher complexity bounds for the similarity learning problem. We give a different formulation of the objective function for learning than [2], which is convex for certain hypothesis classes \mathcal{H} . We then give explicit bounds for both linear and kernel solutions to the optimization problem. We find that, in general, model complexity grows polynomially with the radius of the space. We note, however, that our approach could be used for neural networks as well, since the Rademacher complexity is general. However, we do not yet have the statistical tools to give meaningful upper-bounds on the Rademacher complexity for non-convex hypothesis classes. Therefore, while we will stick with kernel methods in this paper, our research does not preclude the use of more complex hypothesis classes.

1.1 Preliminaries

One simple approach for solving the similarity learning problem is to formulate it as a regression problem. In this case, however, we are not interested in the values of the output space. Instead, we only care about the distance between those points. Therefore, we might update the weights of the regression with respect to the following optimization problem:

$$\operatorname{argmin}_{\theta} \frac{1}{m^2} \mathbf{1}^\top ((\hat{\mathbf{D}}_\theta - \mathbf{D}) \otimes (\hat{\mathbf{D}}_\theta - \mathbf{D})) \mathbf{1}$$

Where \otimes denotes the Hadamard product and $\hat{\mathbf{D}}_\theta$ is the distance matrix between the values computed with our learned parameterized mapping h . In the case of a kernel mapping, we can regularize with the weight matrix (see Corollaries 4.1 and 4.2 for a theoretical justification).

$$\operatorname{argmin}_{\theta} \frac{1}{m^2} \mathbf{1}^\top ((\hat{\mathbf{D}}_\theta - \mathbf{D}) \otimes (\hat{\mathbf{D}}_\theta - \mathbf{D})) \mathbf{1} + \lambda \|w\|_2$$

where λ is the regularization parameter. Also note that, in general, we don't care about the dimensionality of the transformed space \mathcal{Y} , and so a simple way of regularizing would be to lower its dimensionality. Indeed, this is an argument for embedding dimensionality reduction into the learning problem. However, as we will see in Corollary 4.2, the model complexity does not depend on the dimensionality of the data.

As with most learning problems, we are not interested in minimizing the empirical error of a hypothesis $\hat{R}(h)$. Thus, we wish to prove probabilistic bounds on the generalization error $R(h)$ for some hypothesis $h \in \mathcal{H}$. Here, we will assume a fixed, but unknown distance function d . We further assume

that data points $x \in \mathcal{X}$ are drawn i.i.d. according to a fixed, but unknown distribution D .

Generalization error is typically given as a function of the empirical error $\hat{R}(h)$ plus some complexity of the corresponding hypothesis class. A commonly used complexity is the Rademacher complexity $\mathfrak{R}_m(\mathcal{H})$, which is known to have particularly nice, data-dependent properties [1]. Formally, the Rademacher complexity is defined as follows:

Definition 1 (Rademacher complexity [8]). *Let $S = (z_1, \dots, z_m)$ be a fixed sample drawn according to some fixed, but arbitrary distribution D . Then for any sample of size $m > 1$, the Rademacher complexity of a family of functions \mathcal{G} is defined as follows:*

$$\mathfrak{R}_m(\mathcal{G}) = \mathbb{E}_{S, \sigma} \left[\sup_{g \in \mathcal{G}} \frac{1}{m} \sum_{i=1}^m \sigma_i g(z_i) \right]$$

Where $\sigma = (\sigma_1, \dots, \sigma_m)$ with σ_i being a Rademacher random variable, taking values uniformly in $\{-1, +1\}$.

Indeed, one can think of the Rademacher complexity as the ability of a hypothesis class to correlate with random noise [8]. Unfortunately, $\mathfrak{R}_m(\mathcal{H})$ is not typically accessible. However, we can usually get around this problem by deriving upper bounds, which gives us an explicit bound on the generalization error, which we will do in the next section.

2 Theoretical analysis

In this section we give a Rademacher complexity bound for the similarity learning problem. As previously discussed, we wish to minimize the following loss $g = \frac{1}{m^2} \mathbf{1}^\top ((\hat{\mathbf{D}} - \mathbf{D}) \otimes (\hat{\mathbf{D}} - \mathbf{D})) \mathbf{1}$, which is the squared error between the predicted distance matrix $\hat{\mathbf{D}}$ and the true distance matrix \mathbf{D} .

We wish to bound the generalization error for some hypothesis h in a hypothesis class \mathcal{H} . To do this, we will need to make use of the following lemma:

Lemma 2. *Let S and S' be two samples that differ by a single point. We will define $\varphi(S) = \sup_{g \in \mathcal{G}} \mathbb{E}[g] - \hat{\mathbb{E}}_S[g]$. Furthermore, assume that $(\|h(x_i) - h(x_j)\|_2 - \mathbf{D}_{ij})^2 \leq M^2$. Then,*

$$|\varphi(S) - \varphi(S')| \leq \frac{2M}{m}$$

Proof.

$$\begin{aligned}
|\varphi(S') - \varphi(S)| &= \sup_{g \in \mathcal{H}} |\hat{\mathbb{E}}_{S'}[g] - \hat{\mathbb{E}}_S[g]| \\
&= \sup_{g \in \mathcal{H}} \frac{\mathbf{1}^\top ((\hat{\mathbf{D}} - \mathbf{D}) \otimes (\hat{\mathbf{D}} - \mathbf{D})) \mathbf{1} - \mathbf{1}^\top ((\hat{\mathbf{D}}' - \mathbf{D}') \otimes (\hat{\mathbf{D}}' - \mathbf{D}')) \mathbf{1}}{m^2} \leq \frac{2M}{m},
\end{aligned}$$

where we use the fact that changing a single point will change at most $2m$ distances by at most M , by assumption. \square

Lemma 3. Define $\varphi(S)$ as above and let \mathcal{G} be the class of losses over the hypothesis class \mathcal{H} . Then,

$$\mathbb{E}_S[\varphi(S)] \leq 2\mathfrak{R}_m(\mathcal{G})$$

Proof.

$$\begin{aligned}
\mathbb{E}_S[\varphi(S)] &= \mathbb{E}_S \left[\sup_{g \in \mathcal{H}} \mathbb{E}[g] - \hat{E}_S[g] \right] \\
&= \mathbb{E}_S \left[\sup_{g \in \mathcal{H}} \mathbb{E}_{S'} [\hat{\mathbb{E}}_{S'}[g] - \hat{\mathbb{E}}_S[g]] \right] \\
&\leq \mathbb{E}_{S, S'} \left[\sup_{g \in \mathcal{H}} \hat{\mathbb{E}}_{S'}[g] - \hat{\mathbb{E}}_S[g] \right] \\
&= \mathbb{E}_{S, S'} \left[\sup_{g \in \mathcal{H}} \frac{1}{m^2} \mathbf{1}^\top ((\hat{\mathbf{D}}' - \mathbf{D}') \otimes (\hat{\mathbf{D}}' - \mathbf{D}')) \mathbf{1} - \mathbf{1}^\top ((\hat{\mathbf{D}} - \mathbf{D}) \otimes (\hat{\mathbf{D}} - \mathbf{D})) \mathbf{1} \right] \\
&= \mathbb{E}_{S, S'} \left[\sup_{h \in \mathcal{H}} \frac{1}{m^2} \sum_{ij} \left((\|h(x'_i) - h(x'_j)\|_2 - \mathbf{D}'_{ij})^2 - (\|h(x_i) - h(x_j)\|_2 - \mathbf{D}_{ij})^2 \right) \right] \\
&= \mathbb{E}_{\sigma, S, S'} \left[\sup_{h \in \mathcal{H}} \frac{1}{m^2} \sum_{ij} \sigma_{ij} \left((\|h(x'_i) - h(x'_j)\|_2 - \mathbf{D}'_{ij})^2 - (\|h(x_i) - h(x_j)\|_2 - \mathbf{D}_{ij})^2 \right) \right] \\
&\leq \mathbb{E}_{\sigma, S'} \left[\sup_{h \in \mathcal{H}} \frac{1}{m^2} \sum_{ij} \sigma_{ij} (\|h(x'_i) - h(x'_j)\|_2 - \mathbf{D}'_{ij})^2 \right] \\
&\quad + \mathbb{E}_{\sigma, S} \left[\sup_{h \in \mathcal{H}} \frac{1}{m^2} \sum_{ij} \sigma_{ij} (\|h(x_i) - h(x_j)\|_2 - \mathbf{D}_{ij})^2 \right] \\
&= 2 \mathbb{E}_{\sigma, S} \left[\sup_{h \in \mathcal{H}} \frac{1}{m^2} \sum_{ij} \sigma_{ij} (\|h(x_i) - h(x_j)\|_2 - \mathbf{D}_{ij})^2 \right] \\
&= 2\mathfrak{R}_m(\mathcal{G})
\end{aligned}$$

where the second equality holds by the i.i.d. assumption, the first inequality holds by Jensen's inequality, the second inequality holds by the sub-additivity of the supremum function, and we use the fact that $\mathbb{E}[\sigma_i] = \mathbb{E}[-\sigma_i]$. \square

Theorem 4. *Let $S = (x_1, x_2, \dots, x_m)$ be a sample of data drawn i.i.d. from an arbitrary, but fixed distribution D . Let $\mathcal{H} = \{h : (\|h(x_i) - h(x_j)\|_2 - \mathbf{D}_{ij})^2 \leq M^2\}$ and \mathcal{G} be the class of losses over \mathcal{H} . Then, with probability at least $1 - \delta$, for all $h \in \mathcal{H}$*

$$R(h) \leq \hat{R}(h) + 2\mathfrak{R}_m(\mathcal{G}) + M\sqrt{\frac{2\log \frac{1}{\delta}}{m}}$$

Proof. We begin by using McDiarmid's inequality and Lemma 2 to get the following upper bound:

With probability at least $1 - \delta$:

$$|\varphi(S) - \mathbb{E}_S[\varphi(S)]| \leq M\sqrt{\frac{2\log \frac{1}{\delta}}{m}}$$

The bound follows immediately by Lemma 3. \square

Theorem 4 gives us a Rademacher complexity bound on the generalization error for the distances between m points. Notably, the bound depends on M , which is not known in general. However, if the domain of the output space \mathcal{Y} is bounded, then M can be trivially calculated. Below we give an upper bound on the Rademacher complexity for the case where \mathcal{H} is linear.

Corollary 4.1. *Let $\mathcal{H} = \{x \mapsto w \cdot x : \|w\| \leq \Lambda\}$, $\|x\| \leq r$, and $\mathbf{D}_{ij} \leq \beta$. Then, with probability at least $1 - \delta$, for all $h \in \mathcal{H}$:*

$$R(h) \leq \hat{R}(h) + \frac{2M^2}{m} + M\sqrt{\frac{2\log \frac{1}{\delta}}{m}}$$

where $M = \Lambda \max\{2r, \beta\}$.

Proof.

$$\begin{aligned}
\mathfrak{R}_m(\mathcal{G}) &= \mathbb{E}_\sigma \left[\sup_{h \in \mathcal{H}} \frac{1}{m^2} \sum_{ij} \sigma_{ij} (\|h(x_i) - h(x_j)\|_2 - \mathbf{D}_{ij})^2 \right] \\
&= \frac{1}{m^2} \mathbb{E}_\sigma \left[\sup_{\|w\| \leq \Lambda} \sum_{ij} \sigma_{ij} (\|w \cdot x_i - w \cdot x_j\|_2 - \mathbf{D}_{ij})^2 \right] \\
&\leq \frac{\Lambda^2}{m^2} \mathbb{E}_\sigma \left[\sum_{ij} \sigma_{ij} (\|x_i - x_j\|_2 - \mathbf{D}_{ij})^2 \right] \\
&\leq \frac{\Lambda^2}{m^2} \left[\mathbb{E}_\sigma \left[\sum_{ij} \sigma_{ij} (\|x_i - x_j\| - \mathbf{D}_{ij})^2 \right]^2 \right]^{\frac{1}{2}} \\
&= \frac{\Lambda^2}{m^2} \left[\mathbb{E}_\sigma \left[\sum_{ijkl} \sigma_{ij} \sigma_{kl} (\|x_i - x_j\| - \mathbf{D}_{ij})^2 (\|x_k - x_l\| - \mathbf{D}_{kl})^2 \right] \right]^{\frac{1}{2}} \\
&= \frac{\Lambda^2}{m^2} \left[\sum_{ij} (\|x_i - x_j\| - \mathbf{D}_{ij})^4 \right]^{\frac{1}{2}} \\
&\leq \frac{\Lambda^2}{m^2} \left[m^2 \max\{2r, \beta\}^4 \right]^{\frac{1}{2}} = \frac{\Lambda^2}{m} \max\{2r, \beta\}^2
\end{aligned}$$

Where the first inequality holds by the Cauchy-Schwarz inequality and supremum over $\|w\|$. The second inequality holds by Jensen's inequality. The expectation over σ is eliminated because $\mathbb{E}_\sigma[\sigma_i \sigma_j] = 0$ whenever $i \neq j$. The final inequality holds by assumption and the non-negativity of distances. \square

Thus, we can guarantee that the squared distances will be close as long as the radius of the input and output distributions are small relative to the number of samples. Note that this bound also implies that the complexity grows quadratically with the norm of w . More importantly, this bound can be computed directly from the data. A similar bound holds for the kernel formulation.

Corollary 4.2. *Let $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a PDS kernel and let $\phi : \mathcal{X} \rightarrow \mathbb{H}$ be a feature mapping associated to K . Let $S \subseteq \{x : K(x, x) \leq q^2\}$. Let $\mathcal{H} = \{x \mapsto w \cdot \phi(x) : \|w\|_{\mathbb{H}} \leq \Lambda\}$. Then*

$$\mathfrak{R}_m(\mathcal{G}) \leq \frac{\Lambda^2}{m} \max\{\sqrt{2}q, \beta\}^2$$

Proof.

$$\begin{aligned}
\mathfrak{R}_m(\mathcal{G}) &= \frac{1}{m^2} \mathbb{E}_\sigma \left[\sup_{\|w\| \leq \Lambda} \sum_{ij} \sigma_{ij} (\|w \cdot \phi(x_i) - w \cdot \phi(x_j)\|_{\mathbb{H}} - \mathbf{D}_{ij})^2 \right] \\
&\leq \frac{\Lambda^2}{m^2} \mathbb{E}_\sigma \left[\sum_{ij} \sigma_{ij} (\|\phi(x_i) - \phi(x_j)\|_{\mathbb{H}} - \mathbf{D}_{ij})^2 \right] \\
&\leq \frac{\Lambda^2}{m^2} \left[\sum_{ij} (\|\phi(x_i) - \phi(x_j)\|_{\mathbb{H}} - \mathbf{D}_{ij})^4 \right]^{\frac{1}{2}} \\
&= \frac{\Lambda^2}{m^2} \left[\sum_{ij} \left(\sqrt{\|\phi(x_i)\|_{\mathbb{H}}^2 + \|\phi(x_j)\|_{\mathbb{H}}^2 - 2\langle \phi(x_i) | \phi(x_j) \rangle_{\mathbb{H}}} - \mathbf{D}_{ij} \right)^4 \right]^{\frac{1}{2}} \\
&\leq \frac{\Lambda^2}{m^2} \left[\sum_{ij} \left(\sqrt{K(x_i, x_i) + K(x_j, x_j)} - \mathbf{D}_{ij} \right)^4 \right]^{\frac{1}{2}} \leq \frac{\Lambda^2}{m} \max\{\sqrt{2}q, \beta\}^2
\end{aligned}$$

where the first and second inequalities follow by the same logic in Corollary 4.1. The second equality is a simple identity. The third inequality holds because K is PDS and the final inequality holds by assumption. \square

As we can see, the Rademacher complexity depends on the maximum of the square root of the diameter of the data in the Hilbert space defined by the kernel K and the radius of the distance space. In some cases, this bound is actually tighter than the one in Corollary 4.1. In particular, when K is the RBF kernel, $q = 1$. Thus, in this case, the bound is independent of the dimensionality of \mathcal{X} .

3 Conclusion

We presented a theoretical analysis of the similarity learning problem and gave probabilistic learning guarantees. By formulating the problem as a kernel regression from \mathcal{X} to \mathcal{Y} , we were able to give explicit, data dependent bounds, based on the Rademacher complexity of the hypothesis class \mathcal{H} . We showed that, in the linear case, model complexity depended on the radius of the data in either \mathcal{X} or \mathcal{Y} . Notably, we showed that when learning with kernels, the bound is independent of the dimensionality of \mathcal{X} .

References

- [1] Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- [2] Qiong Cao, Zheng-Chu Guo, and Yiming Ying. Generalization bounds for metric and similarity learning. *Machine Learning*, 102(1):115–132, 2016.
- [3] James J DiCarlo and David D Cox. Untangling invariant object recognition. *Trends in cognitive sciences*, 11(8):333–341, 2007.
- [4] Chou P Hung, Gabriel Kreiman, Tomaso Poggio, and James J DiCarlo. Fast readout of object identity from macaque inferior temporal cortex. *Science*, 310(5749):863–866, 2005.
- [5] Roozbeh Kiani, Hossein Esteky, Koorosh Mirpour, and Keiji Tanaka. Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *Journal of neurophysiology*, 97(6):4296–4309, 2007.
- [6] Nikolaus Kriegeskorte, Marieke Mur, N Kriegeskorte, and G Kreiman. Representational similarity analysis of object population codes in humans, monkeys, and models. *Visual Population Codes: Towards a Common Multivariate Framework for Cell Recording and Functional Imaging*, 2012.
- [7] Nikolaus Kriegeskorte, Marieke Mur, Douglas A Ruff, Roozbeh Kiani, Jerzy Bodurka, Hossein Esteky, Keiji Tanaka, and Peter A Bandettini. Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, 60(6):1126–1141, 2008.
- [8] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2012.